Electron. Commun. Probab. **21** (2016), no. 82, 1–12.
DOI: 10.1214/16-ECP29
ISSN: 1083-589X

ELECTRONIC COMMUNICATIONS in PROBABILITY

arXiv:1607.07914v2 [math.PR] 27 Dec 2016

The critical density for the frog model is the degree of the tree*

Tobias Johnson† Matthew Junge ‡

Abstract

The frog model on the rooted d-ary tree changes from transient to recurrent as the number of frogs per site is increased. We prove that the location of this transition is on the same order as the degree of the tree.

Keywords: frog model; transience; recurrence; phase transition; trees; stochastic dominance.
AMS MSC 2010: 60K35; 60J80; 60J10.
Submitted to ECP on July 28, 2016, final version accepted on November 15, 2016.

1 Introduction

The *frog model* in its most general form is an interacting particle system taking place on a countable collection of vertices, typically a graph, with a designated vertex $\varnothing$ that we call the root. The process is given by a pair (η, S) of particle counts and paths. Initially, there is one active particle at the root, and $\eta(v)$ sleeping particles at each nonroot vertex v. When activated, the ith particle starting at vertex v moves according to the path $S_\bullet(v,i)$, with $S_0(v,i)$ assumed to be v. Whenever an active particle moves to a new vertex, it activates all sleeping particles there. We continue the common practice of referring to the particles as frogs.

Let $\mathbb{T}_d$ denote the infinite d-ary tree rooted at $\varnothing$; that is, $\mathbb{T}_d$ is the tree in which $\varnothing$ has degree d and all other vertices have degree $d+1$. Our topic is the frog model on $\mathbb{T}_d$ where $(\eta(v))_{v\neq\varnothing}$ is an i.i.d.-$\mathrm{Poi}(\mu)$ collection of random variables, $(S_\bullet(v,i))_{v,i}$ is a collection of independent simple random walk paths, and the two collections are independent. We abbreviate this as the frog model on $\mathbb{T}_d$ with i.i.d.-$\mathrm{Poi}(\mu)$ frogs per site.

We call a realization of the frog model *transient* if the root is visited finitely many times by frogs and *recurrent* if it is visited infinitely often. Recurrence versus transience is perhaps the most basic question for the frog model. It has been studied on $\mathbb{Z}^d$ under a variety of initial conditions and frog paths [TW99, Pop01, GS09, DP14, KZ16, Ros16]. In [HJJ16b, HJJ16a, JJ16], we address the question on d-ary trees. For more background material on the frog model, see the introduction to [JJ16].

We show in [HJJ16a] that the frog model on a d-ary tree with i.i.d.-$\mathrm{Poi}(\mu)$ frogs per vertex undergoes a phase transition between transience and recurrence as μ grows. In more detail, for each $d \geq 2$ there exists a critical value $\mu_c(d)$ such that the model is almost surely transient for $\mu < \mu_c(d)$ and is almost surely recurrent for $\mu > \mu_c(d)$. The proof in

*The first author was partially supported by NSF grant DMS-1401479.
†New York University. E-mail: tobias.johnson@nyu.edu
‡Duke University. E-mail: jungem@math.duke.edu

[HJJ16a] shows that for some constants $C, C' > 0$, we have $Cd \leq \mu_c(d) \leq C' d \log d$. In this paper, we sharpen this result by removing the $\log d$ factor from the upper bound, determining the order of $\mu_c(d)$ up to constant factors.

Theorem 1. For all sufficiently large d, it holds that $.24d \leq \mu_c(d) \leq 2.28d$.

Theorem 1 can be extended to initial distributions other than Poisson using the results of [JJ16]. For example, the frog model on a d-ary tree with deterministically many frogs per site has a critical threshold of the same order; see [JJ16, Corollary 4] for details.

An accurate description of the transition threshold on trees is especially relevant given that the frog model on $\mathbb{Z}^d$ behaves rather differently. A transition still occurs, but it does so at a decaying density of frogs (see [Pop01, Theorem 1.1]). A natural next step would be to investigate this problem for irregular trees. Perhaps the phase transition is on the same order as the branching number of the tree (see [LP16, Section 1.2])? This is completely speculative, as the mere existence of a phase transition for the frog model is unknown even on a Galton–Watson tree.

Our other goal for this paper is to present as simply as possible the argument for existence of a recurrence phase on trees. While the transience phase is fairly easy to establish (see [HJJ16a, Proposition 15]), the recurrence phase is more difficult. Indeed, the question of recurrence on $\mathbb{T}_2$ with one sleeping frog per site was first posed in [AMP02] and was only recently answered in the affirmative [HJJ16b]. It remains open to determine if the one-per-site frog model is recurrent on $\mathbb{T}_3$ and $\mathbb{T}_4$. Taking advantage of some technical improvements we have made since [HJJ16b, HJJ16a], we give a streamlined proof of recurrence on $\mathbb{T}_2$ with i.i.d.-$\mathrm{Poi}(\mu)$ frogs per site (see Proposition 3). Using [JJ16, Corollary 3], this result also implies recurrence on $\mathbb{T}_2$ with two sleeping frogs per site. This is weaker than the result in [HJJ16b], but the proof is much simpler.

Ideas of the proof

As in [HJJ16b, HJJ16a], our proof of recurrence is based on recursion and bootstrapping. To set this up, we first show that it is enough to establish recurrence for a frog model whose paths are stopped non-backtracking walks, which we call the *self-similar frog model*. Let V be the number of visits to the root in this process. A self-similarity yields a relation between V and a collection of independent copies of V. Such relations are called recursive distributional equations (see [AB05] for further discussion).

In the bootstrap part of the argument, we assume that V is stochastically larger than $\mathrm{Poi}(\lambda)$ for some $\lambda \geq 0$. We then analyze the recursive distributional equation to show that V is in fact stochastically larger than $\mathrm{Poi}(\lambda + \epsilon)$. Iterating this argument starting at $\lambda = 0$, we show that V is larger than $\mathrm{Poi}(\epsilon)$, then larger than $\mathrm{Poi}(2\epsilon)$, and so on, with the conclusion that $V = \infty$ a.s. Here and in [HJJ16a], this argument uses the standard stochastic order, while in [HJJ16b] it uses a more exotic stochastic order.

The bootstrap phase of our argument is more elaborate than in [HJJ16a], where we could only establish $\mu_c(d) = O(d \log d)$. We obtain a better upper bound in this paper because the version of the self-similar frog model here better approximates the actual frog model. Put more simply, we are able to capture the contributions of more frogs. The simplification that allows us to handle the extra complexity is our use of [MSH03, Theorem 3.1(b)], a simple criterion for determining when a Poisson distribution stochastically dominates a Poisson mixture. This allows us to avoid the more difficult coupling argument used in [HJJ16a].

2 A criterion for stochastic dominance

Given two probability measures ν and ν' on the extended real numbers, we say that ν is *stochastically smaller* than ν' if

$$\nu\big((t,\infty]\big) \le \nu'\big((t,\infty]\big)$$

for all $t \in \mathbb{R}$. We denote this relationship by $\nu \preceq \nu'$. If $X \sim \nu$ and $Y \sim \nu'$, we also write $X \preceq Y$, $X \preceq \nu'$, and $\nu \preceq Y$ all to mean the same thing. An alternate characterization of stochastic dominance is that $X \preceq Y$ if and only if there exists a coupling of X and Y such that $X \le Y$ a.s.

Theorem 2 provides a necessary and sufficient condition for a Poisson mixture to dominate a Poisson distribution, which we will need in Section 3. We reproduce the proof given in [MSH03] for our readers' convenience. See also [Yu09] for a more general result.

Lemma 2.1 ([MSH03, Lemma 3.1(b)])**.** *For any positive integer n, the function*

$$h_n(x) = x \sum_{k=0}^{n} \frac{(-\log x)^k}{k!}$$

is increasing and concave on $(0,1]$.

Proof. We compute

$$h_n'(x) = \sum_{k=0}^{n} \frac{(-\log x)^k}{k!} - \sum_{k=1}^{n} \frac{(-\log x)^{k-1}}{(k-1)!} = \frac{(-\log x)^n}{n!},$$

which is positive and decreasing on $(0,1]$, showing that $h_n(x)$ is increasing and concave. □

Theorem 2 ([MSH03, Theorem 3.1(b)])**.** Let $X \sim \mathrm{Poi}(\lambda)$, and let $Y \sim \mathrm{Poi}(U)$ for some nonnegative random variable U. Then the following are equivalent:

(i) $X \preceq Y$,

(ii) $\mathbf{P}[X=0] \ge \mathbf{P}[Y=0]$, and

(iii) $\lambda \le -\log \mathbf{E}e^{-U}$.

Proof. Conditions (ii) and (iii) are just restatements of each other, since $\mathbf{P}[X=0] = e^{-\lambda}$ and $\mathbf{P}[Y=0] = \mathbf{E}e^{-U}$. Condition (i) implies (ii) by the definition of stochastic dominance. It remains to prove that (iii) implies (i). It suffices to show that $\mathbf{P}[Y \le n] \le \mathbf{P}[X \le n]$ for all nonnegative integers n. We compute

$$\mathbf{P}[Y \le n] = \sum_{k=0}^{n} \mathbf{E}\left[\frac{e^{-U}U^k}{k!}\right] = \mathbf{E}\left[\zeta \sum_{k=0}^{n} \frac{(-\log \zeta)^k}{k!}\right] = \mathbf{E}h_n(\zeta),$$

where $\zeta = e^{-U}$. Our assumption is that $\mathbf{E}\zeta = \mathbf{E}e^{-U} \le e^{-\lambda}$. By Lemma 2.1, the function $h_n(x)$ is increasing and concave on $(0,1]$, where ζ takes values. Thus

$$\mathbf{E}h_n(\zeta) \le h_n(\mathbf{E}\zeta) \le h_n\big(e^{-\lambda}\big) = \sum_{k=0}^{n} \frac{e^{-\lambda}\lambda^k}{k!} = \mathbf{P}[X \le n],$$

where we use that h_n is concave to apply Jensen's inequality in the first step, and we use that h_n is increasing in the second step. □

3 Critical parameters for d-ary trees

Our argument breaks down into two parts described in the introduction. In Section 3.1, we define the self-similar frog model and define V as the number of visits to the root in this process. Then, we determine a recursive distributional equation satisfied by the law of V (see Lemma 3.5). The ideas in this section can be found in [HJJ16b, HJJ16a], but they take some work to extract in the form we need. Though we do our best to avoid duplicating material, when in doubt we have opted for comprehensibility over efficiency.

Next comes the bootstrap portion of the argument. In Section 3.2, we use the set-up of Section 3.1 to carry this out in case of the binary tree, giving a short proof of the existence of a recurrence phase. In Section 3.3 we give a more complex version of this argument proving recurrence on the d-ary tree for $\mu = \Omega(d)$.

3.1 The bootstrapping set-up

3.1.1 The self-similar frog model

Many basic features of the frog model depend only on the range of each frog. This yields rather nice abelian and monotonicity properties. For example, the total number of visits to the root is unaffected by the order frogs wake up in and the rate they reveal vertices in their ranges. Also, trimming the range of frogs can only reduce the number of visits to the root. Applying this observation in combination with the coupling characterization of stochastic dominance, we note the following fact. Let $r(\eta, S)$ be the number of visits to the root in the frog model (η, S).

Fact 3.1. *Consider a collection of frog paths $S = \big(S_\bullet(v,i)\big)_{v\in G, i\geq 1}$ on a graph G. Suppose that another collection of paths $\tilde{S}$ can be coupled with S such that for all i and v, the range of $\tilde{S}_\bullet(v,i)$ is a subset of the range of $S_\bullet(v,i)$. Then $r(\eta,\tilde{S}) \preceq r(\eta, S)$.*

From now on, let $S = (S_\bullet(v,i),\, v \in \mathbb{T}_d, i \geq 1)$ denote a collection of independent simple random walks with $S_\bullet(v,i)$ started at v, and let the components of $\eta = (\eta(v))_{v\in\mathbb{T}_d}$ be i.i.d.-$\mathrm{Poi}(\mu)$, independent of S. The first step in studying the frog model (η, S) will be to replace S by a collection of paths T to obtain (η, T), which we call the *self-similar frog model* in reference to a useful property described in Fact 3.3.

We define T in two steps. First, let $S' = (S'_\bullet(v,i),\, v \in \mathbb{T}_d, i \geq 1)$ denote a collection of independent random non-backtracking walks stopped at $\varnothing$. In more detail, call a random walk a *simple random non-backtracking walk* on an arbitrary graph if it chooses from its neighbors uniformly for its first step, and then in all subsequent steps it chooses uniformly from its current neighbors except the one it just arrived from. We define $S'_\bullet(v,i)$ to be a simple non-backtracking random walk stopped on arrival at $\varnothing$. The walks $S'_\bullet(v,i)$ and $S_\bullet(v,i)$ can be coupled so that the range of the first is a subset of the range of the second by making $S'_\bullet(v,i)$ a stopped, loop-erased version of $S_\bullet(v,i)$. This is proved in detail in [HJJ16b, Proposition 7].

Now we construct T as a modification of S'. Each path $T_\bullet(v,i)$ will be a stopped version of $S'_\bullet(v,i)$. Let v be a nonroot vertex in $\mathbb{T}_d$ with parent u. Suppose that v is visited in the frog model (S', η) for the first time at time j, necessarily by one or more frogs moving from u to v. Select one of these visiting frogs arbitrarily, and stop all of the other ones. (Observe that it is irrelevant which frog is allowed to continue, so long as one views frogs as indistinguishable.) If any frogs move from u to v at subsequent times, stop them at v as well. Do this for all vertices $v \in \mathbb{T}_d$, and let T be the resulting collection of stopped walks. As the range of each $T_\bullet(v,i)$ is a subset of the range of $S'_\bullet(v,i)$, the following fact (also noted in [HJJ16b, Proposition 7]) follows:

Fact 3.2. *There is a coupling of S and T so that the range of each $T_\bullet(v,i)$ is a subset of the range of $S_\bullet(v,i)$.*

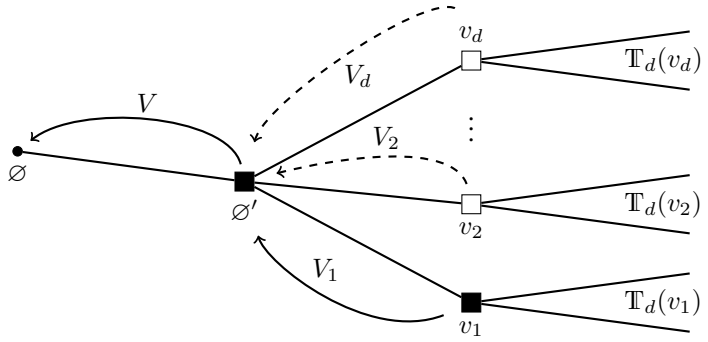

Figure 1: In the self-similar frog model on $\mathbb{T}_d$, the initial frog moves from $\varnothing$ to $\varnothing'$ to v_1 and then continues down the tree. The random variables V and V_1, counting the number of frogs moving from $\varnothing'$ to $\varnothing$ and from v_1 to $\varnothing'$, respectively, both have distribution ν (see Lemma 3.5). For $i \in \{2, \dots, d\}$, the distribution of V_i conditional on a frog entering the subtree $\mathbb{T}_d(v_i)$ is also ν.

By Facts 3.1 and 3.2, we have $r(\eta, T) \preceq r(\eta, S)$. We will now work exclusively with the self-similar frog model, (η, T), and prove recurrence for it with sufficiently large μ. Unlike all other frog models considered in this paper, the frog paths T are not independent of each other nor of η, because one frog's motion in (η, T) can cause another frog to be stopped. This is the only form of dependence, however, and frogs that have not been stopped move independently of each other. So, it is not a serious obstacle.

Let $V = r(\eta, T)$. Next, we discuss a self-similarity property of the model and its consequences for V. For any vertex $v \in \mathbb{T}_d$, let $\mathbb{T}_d(v)$ denote the subtree made up of v and its descendants. We call $\mathbb{T}_d(v)$ *activated* in the self-similar frog model if v is ever visited. Let u be the parent of v. By our construction of T, if $\mathbb{T}_d(v)$ is activated, then there is a unique frog that moves from u to v, entering $\mathbb{T}_d(v)$ and then never leaving it. The frog model viewed starting from the time of activation only at vertices $\{u\} \cup \mathbb{T}_d(v)$ then looks identical to the original self-similar frog model viewed on $\{\varnothing\} \cup \mathbb{T}_d(\varnothing')$. This yields the following fact, proved in more detail in [HJJ16b, Proposition 6].

Fact 3.3. *Let V' be the number of frogs that move from v to its parent u in the self-similar frog model. The distribution of V' conditional on $\mathbb{T}_d(v)$ being activated is identical to the distribution of V.*

The following observation shows that once a subtree $\mathbb{T}_d(v)$ is activated, the random variable V' defined in the above fact is independent of the frog model outside of $\mathbb{T}_d(v)$.

Fact 3.4. *Let V' be defined as in Fact 3.3. Conditional on $\mathbb{T}_d(v)$ being activated, V' depends only on the path of the activator and on $\{T_\bullet(w, i), \eta(w) \colon w \in \mathbb{T}_d(v), i \geq 1\}$.*

We will use Facts 3.3 and 3.4 to express V recursively in terms of independent copies of itself, an idea expressed in Figure 1. This relation will be given in terms of an operator we define next.

3.1.2 The operators $\mathcal{B}$ and $\mathcal{U}$

Suppose that the initial frog in the self-similar frog model moves from $\varnothing$ to $\varnothing'$ to v_1. Let $v_2, \dots, v_d$ be the remaining children of $\varnothing'$. Observe that since frogs are stopped at $\varnothing$, no children of $\varnothing$ other than $\varnothing'$ are ever visited. The idea of this section is to view the self-similar frog model only at the vertices mentioned above. If a vertex v_i is visited, we

close our eyes to $\mathbb{T}_d(v_i)$, thinking of this entire subtree as a black box that eventually emits some frogs from v_i back to $\varnothing'$.

Enacting this view, we now define operators $\mathcal{B}$ and $\mathcal{U}$ on probability measures supported on the extended nonnegative integers. Informally, the operator $\mathcal{B}$ corresponds to the number of visits to the root, and $\mathcal{U}$ corresponds to the number of subtrees $v_1, \dots, v_d$ that are activated. Let π be a probability measure on the nonnegative integers. To define $\mathcal{B}\pi$ and $\mathcal{U}\pi$, we consider the following frog model. The example to keep in mind is when π is the law of V, in which case the following description matches up with the black box view of the self-similar frog model described above.

Graph a star graph with center ρ' and leaf vertices $\rho, u_1, \dots, u_d$ (think of these as paralleling $\varnothing'$ and $\varnothing, v_1, \dots, v_d$). The root of the graph is ρ.

Sleeping frog counts all independent, distributed as $\mathrm{Poi}(\mu)$ at ρ' and as π at $u_1, \dots, u_d$. There is one frog at ρ, as is always true at the root vertex.

Paths All frogs have independent paths. The initial frog moves deterministically from ρ to ρ' to u_1 and then remains there. All other frogs, if woken, perform simple random non-backtracking walks from their starting points, stopped on arrival at a leaf vertex.

We then define two quantities:

- $\mathcal{B}\pi$ is the distribution of the number of frogs that terminate at ρ.
- $\mathcal{U}\pi$ is the distribution of the final number of $u_1, \dots, u_d$ that are visited by a frog.

Note that our definition of the initial frog path as deterministic is just for convenience. By symmetry, we would arrive at the same measures $\mathcal{B}\pi$ and $\mathcal{U}\pi$ if it were also defined as a stopped simple random non-backtracking walk.

We mention that $\mathcal{B}$ is closely related to the operators $\mathcal{A}$ defined in [HJJ16b] and [HJJ16a], but differs from both of them. The operator $\mathcal{A}$ in [HJJ16b] is the same as $\mathcal{B}$ in the $d = 2$ case if the initial distribution at ρ' in the definition of $\mathcal{B}$ is changed from $\mathrm{Poi}(\mu)$ to δ_1 (except that $\mathcal{A}$ acts on probability generating functions rather than distributions). The operator $\mathcal{A}$ in [HJJ16a] is the same as $\mathcal{B}$ in the $d = 2$ case. For $d \geq 3$, the two operators differ in that in the frog model defining $\mathcal{A}$, frogs initially at $v_2, \dots, v_d$ do not wake other frogs.

Now we relate this system back to the frog model.

Lemma 3.5. *Let ν be the law of $V = r(\eta, T)$, the number of visits to the root in the self-similar frog model on $\mathbb{T}_d$. It holds that $\mathcal{B}\nu = \nu$.*

Proof. Essentially, the frog model on the star graph exactly matches the black box view of the self-similar frog model described at the beginning of Section 3.1.2, and the result then follows from Facts 3.3 and 3.4. To make this more formal, we couple the two frog models. We take full advantage of the abelian properties of the frog model by viewing the frogs' motions in a convenient order.

Consider the frog model used to define $\mathcal{B}\nu$ as well as the self-similar frog model. We can couple the initial number of frogs on ρ' to be the same as on $\varnothing'$, and we can couple the first (and only) step of each frog at ρ' with the first step of the corresponding frog at $\varnothing'$.

Let V_i be the number of frogs that ever move from v_i to $\varnothing'$ in the self-similar model, and let U_i be the number of frogs initially at u_i in the star graph model. Noting that $\mathbb{T}_d(v_1)$ is activated by the initial frog, $V_1 \sim \nu$ by Fact 3.3. By Fact 3.4, V_1 is independent of all that we have coupled so far (that is, the number and first steps of frogs initially at $\varnothing'$). The random variable U_1 is also independent of all we have coupled so far and is

distributed identically to V_1. We can therefore couple U_1 and V_1 to be equal. Next, we couple the second (and final) step of each frog at u_1 with the step of the corresponding frog counted by V_1 after it moves from v_1 to $\varnothing'$.

Let $\mathcal{V}_1$ consist of the indices $i \in \{2, \dots, d\}$ such that u_i has been visited so far. By the construction of our coupling, we can also describe $\mathcal{V}_1$ as the set of $i \in \{2, \dots, d\}$ such that $\mathbb{T}_d(v_i)$ has been activated so far. Furthermore, identically many frogs have returned so far to ρ as to $\varnothing$. By Facts 3.3 and 3.4, conditional on the information so far, the random variables $(V_i,\, i \in \mathcal{V}_1)$ are i.i.d.-ν and are independent of the information so far, as are the random variables $(U_1,\, i \in \mathcal{V}_1)$. We can therefore couple these two random vectors to be equal. We then couple the paths of the frogs at these vertices to match up as we did with the frogs at u_1 and v_1.

As above, a vertex u_i is visited for the first time in this second round if and only if $\mathbb{T}_d(v_i)$ is visited for the first time in this second round. Let $\mathcal{V}_2$ be the set of such i. We can repeat the coupling argument of the previous paragraph, maintaining identical numbers of frogs terminating at ρ as at $\varnothing$, until we get an empty $\mathcal{V}_j$ and have counted all returns to ρ and $\varnothing$. Thus, under this coupling, the number of frogs terminating at ρ in the star graph model is the same as the number of frogs terminating at $\varnothing$ in the self-similar model. The first of these counts has distribution $\mathcal{B}\nu$, while the second has distribution ν, showing that the two are equal. □

The next lemma is similar to [HJJ16b, Lemma 10] and [HJJ16a, Lemma 10].

Lemma 3.6. *If $\pi_1 \preceq \pi_2$, then $\mathcal{B}\pi_1 \preceq \mathcal{B}\pi_2$.*

Proof. This immediately follows from the coupling definition of stochastic dominance. We couple the frog models defining $\mathcal{B}\pi_1$ and $\mathcal{B}\pi_2$ so that the frogs in the former are a subset of the frogs of the latter model, resulting in more visits to the root. □

Just as in [HJJ16a, Lemma 11], the operator $\mathcal{B}$ applied to a Poisson distribution yields a mixture of Poisson distributions. This is a consequence of the following property, known as *Poisson thinning*: Consider a multinomial distribution with $\mathrm{Poi}(\lambda)$ trials and n-types, each having probability p_k. Then the vector of outcomes is distributed as an independent collection of $\mathrm{Poi}(\lambda p_k)$ random variables.

Lemma 3.7. *Let U be a random variable distributed as $\mathcal{U}\,\mathrm{Poi}(\lambda)$.*

$$\mathcal{B}\,\mathrm{Poi}(\lambda) = \mathrm{Poi}\left(\frac{\mu}{d+1} + U\frac{\lambda}{d}\right). \tag{3.1}$$

Proof. In the frog model defining $\mathcal{B}\,\mathrm{Poi}(\lambda)$, the number of frogs at ρ' that move back to ρ is distributed as $\mathrm{Bin}(\mathrm{Poi}(\mu), 1/(d+1))$. By Poisson thinning, this is $\mathrm{Poi}(\mu/(d+1))$. Each visited u_i releases $\mathrm{Poi}(\lambda)$ sleeping frogs. These will take a non-backtracking step back to ρ', then with probability $1/d$ will move to ρ. Thus, each activated u_i sends $\mathrm{Poi}(\lambda/d)$ frogs to ρ. It follows that

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \sim \mathrm{Poi}(\mu/(d+1)) + \textstyle\sum_1^d \mathbf{1}\{u_i \text{ visited}\}\,\mathrm{Poi}(\lambda/d).$$

The above sum is equal to $\sum_1^U \mathrm{Poi}(\lambda/d)$. By Poisson thinning, the $\mathrm{Poi}(\lambda/d)$ terms are independent of U. Applying additivity of Poisson random variables then brings us to the claimed formula. □

3.2 Simplest proof of recurrence on the binary tree

We now break from the main thread to give a short proof that the frog model on the binary tree with Poisson-distributed frogs has a recurrence phase. The idea of the argument is to use Lemma 3.7 to demonstrate that for some $\delta > 0$, it holds for all $\lambda \geq 0$

that $\mathcal{B}\operatorname{Poi}(\lambda) \succeq \operatorname{Poi}(\lambda+\delta)$. Lemmas 3.5 and 3.6 then let us bootstrap out way to the conclusion that that V is stochastically larger than any Poisson distribution, and hence $V=\infty$ a.s.

Proposition 3. The frog model on $\mathbb{T}_2$ with i.i.d.-$\operatorname{Poi}(\mu)$ frogs per site is recurrent for $\mu > 3\log\big((1+\sqrt{5})/2\big) \approx 1.4436$.

Proof. In the $d=2$ case, $\mathcal{B}\operatorname{Poi}(\lambda)$ is a particularly simple mixture:

$$\mathcal{B}\operatorname{Poi}(\lambda) = \begin{cases} \operatorname{Poi}\big(\frac{\mu}{3}+\frac{\lambda}{2}\big) & \text{with probability } \exp\big(-\frac{\mu}{3}-\frac{\lambda}{2}\big), \\ \operatorname{Poi}\big(\frac{\mu}{3}+\lambda\big) & \text{with probability } 1-\exp\big(-\frac{\mu}{3}-\frac{\lambda}{2}\big). \end{cases} \tag{3.2}$$

This follows from Lemma 3.7 once we show that

$$\mathcal{U}\operatorname{Poi}(\lambda) = \begin{cases} 1 & \text{with probability } \exp\big(-\frac{\mu}{3}-\frac{\lambda}{2}\big), \\ 2 & \text{with probability } 1-\exp\big(-\frac{\mu}{3}-\frac{\lambda}{2}\big). \end{cases} \tag{3.3}$$

Recall that $\mathcal{U}\operatorname{Poi}(\lambda)$ is the distribution of the number of vertices out of $\{u_1,u_2\}$ that are visited in the frog model defining $\mathcal{B}$. The vertex u_1 is always visited by the initial frog in this model. Each of the $\operatorname{Poi}(\mu)$ frogs initially at ρ' has a $1/3$ chance of visiting u_2, and each of the $\operatorname{Poi}(\lambda)$ frogs initially at u_1 has a $1/2$ chance of visiting u_2, all independently of each other. By Poisson thinning, the number of these frogs that visit u_2 is distributed as $\operatorname{Poi}(\mu/3+\lambda/2)$, and thus u_2 is visited with probability $1-\exp(-\mu/3-\lambda/2)$. This establishes (3.3) and hence (3.2).

Theorem 2 now instructs us to compute the probability placed on 0 by $\mathcal{B}\operatorname{Poi}(\lambda)$, which is

$$\begin{aligned} -\log\Big[\big(1-e^{-\frac{\mu}{3}-\frac{\lambda}{2}}\big)e^{-\frac{\mu}{3}-\lambda}+e^{-\frac{2\mu}{3}-\lambda}\Big] &= \lambda+\frac{\mu}{3}-\log\Big[1-e^{-\frac{\mu}{3}-\frac{\lambda}{2}}+e^{-\frac{\mu}{3}}\Big] \\ &\geq \lambda+\frac{\mu}{3}-\log\Big[1+e^{-\frac{\mu}{3}}\Big] > \lambda+\delta \end{aligned}$$

for some $\delta>0$ depending on μ but not on λ, under our assumption that $\mu > 3\log\big((1+\sqrt{5})/2\big)$. By Theorem 2,

$$\mathcal{B}\operatorname{Poi}(\lambda) \succeq \operatorname{Poi}(\lambda+\delta) \tag{3.4}$$

for any $\lambda \geq 0$.

Now, we carry out the bootstrap. Recall that ν is the distribution of $r(\eta,T)$, the number of visits to the root in the nonbacktracking frog model on the binary tree. As $\nu \succeq \operatorname{Poi}(0)$, Lemma 3.6 shows that $\mathcal{B}\nu \succeq \mathcal{B}\operatorname{Poi}(0)$, and so $\mathcal{B}\nu \succeq \operatorname{Poi}(\delta)$ by (3.4). But ν is a fixed point of $\mathcal{B}$ by Lemma 3.5, implying that $\nu \succeq \operatorname{Poi}(\delta)$. Repeating this argument of successively applying Lemma 3.6, (3.4), and Lemma 3.5, we show that $\nu \succeq \operatorname{Poi}(2\delta)$, and so on. Thus ν is stochastically larger than all Poisson distributions, which implies $\nu = \delta_\infty$. Finally, recalling that $r(\eta,S)$ is the number of visits to the root in the frog model and that $r(\eta,S) \succeq \nu$ by Facts 3.1 and 3.2, we can conclude that $r(\eta,S)=\infty$ a.s. □

3.3 A more complicated bootstrap

We now give an argument along the same lines as our proof of Proposition 3, proving recurrence for all d under assumptions on μ that are optimal up to constant factors, as shown by the lower bound in Theorem 1. This bound follows from [HJJ16a, Proposition 15], which is proven by coupling the frog model with a transient branching random walk. Our contribution here is the upper bound.

Recall that $\mathcal{U}\operatorname{Poi}(\lambda)$ is the distribution of the number of vertices $u_1, \dots, u_d$ visited in the frog model on the star graph defined in Section 3.1.2. The essential difference between our proofs of recurrence for $\mu = \Omega(d)$ here and for $\mu = \Omega(d \log d)$ in [HJJ16a] is that here we give a better lower bound on $\mathcal{U}\operatorname{Poi}(\lambda)$. For a fixed $\lambda \geq 0$, we define a lower bounding random variable $U' \in \{1, \dots, d\}$ as follows. Consider the frog model used to define $\mathcal{B}\operatorname{Poi}(\lambda)$ and $\mathcal{U}\operatorname{Poi}(\lambda)$, and observe how many of $u_1, \dots, u_d$ are visited by the $\operatorname{Poi}(\mu)$ frogs starting at ρ'. If at least $\lceil d/c \rceil$ of these vertices are visited for a yet to be determined constant c, then arbitrarily choose $\lceil d/c \rceil$ of them and allow the frogs activated there the chance to visit the remaining $d - \lceil d/c \rceil$ vertices. If fewer than $\lceil d/c \rceil$ vertices are visited by the frogs at ρ', then recall that u_1 is guaranteed to be activated by the initial frog, and just use the frogs at u_1 to try to activate the remaining vertices $u_2, \dots, u_d$. We define U' as the number of vertices out of $u_1, \dots, u_d$ activated in the end in this scheme. This is summarized as follows:

Let U_1' be the number of vertices $u_1, \dots, u_d$ visited by the frogs initially at ρ'.

Case 1 $U_1' \geq \lceil d/c \rceil$

Arbitrarily choose $\lceil d/c \rceil$ of the vertices counted by U_1' and denote them by $\mathcal{V} \subseteq \{u_1, \dots, u_d\}$. Let U' be the sum of $\lceil d/c \rceil$ and the number of the remaining $d - \lceil d/c \rceil$ vertices $\{u_1, \dots, u_d\} \setminus \mathcal{V}$ visited by frogs starting in $\mathcal{V}$.

Case 2 $U_1' < \lceil d/c \rceil$

Let U' equal one plus the number of number of vertices $u_2, \dots, u_d$ visited by frogs returning from u_1.

As U' counts only a subset of the full collection of activated vertices, we have $U' \preceq \mathcal{U}\operatorname{Poi}(\lambda)$.

We now sketch the argument in more detail. Throughout the remainder of the paper, we will assume that $\mu = C(d+1)$ with C a yet to be determined positive constant. In Lemma 3.8, we prove that Case 2 occurs with exponentially small probability as d grows. Next, in Lemma 3.9 we give a very explicit definition of a random variable U'' satisfying $U'' \preceq U' \preceq \mathcal{U}\operatorname{Poi}(\lambda)$. In Lemma 3.10, we use this lower bound together with Lemma 3.7 to prove that if $V \succeq \operatorname{Poi}(\lambda)$, then $V \succeq \operatorname{Poi}(\lambda + \delta)$ for some $\delta > 0$. The same iterative argument used in Proposition 3 then implies that $V = \infty$ a.s.

Lemma 3.8. *Recall that U_1' is the number of vertices $u_1, \dots, v_d$ visited by the $\operatorname{Poi}(\mu)$ frogs initially at ρ' in the frog model defining $\mathcal{B}\operatorname{Poi}(\lambda)$ and $\mathcal{U}\operatorname{Poi}(\lambda)$. We have*

$$\mathbf{P}[U_1' < \lceil d/c \rceil] \leq e^{-bd} := p, \tag{3.5}$$

where $b = 2(1 - e^{-C} - \frac{1}{c})^2$.

Proof. It is a consequence of Poisson thinning that out of the $\operatorname{Poi}(\mu)$ frogs starting at ρ', independently $\operatorname{Poi}(\frac{\mu}{d+1}) = \operatorname{Poi}(C)$ move to each leaf $u_1, \dots, u_d$. Thus each vertex has an independent $1 - e^{-C}$ chance of having a frog visit it from the ones starting at ρ', showing that $U_1' \sim \operatorname{Bin}(d, 1 - e^{-C})$.

Hoeffding's inequality tailored to a binomial distribution states that $\mathbf{P}[\operatorname{Bin}(n,p) \leq (p - \epsilon)n] \leq \exp(-2\epsilon^2 n)$ (this follows from [Hoe63, eq. (2.3)]). If we apply the inequality to U_1' with $\epsilon = (1 - e^{-C}) - \frac{1}{c}$, we establish (3.5). □

Lemma 3.9. *Let*

$$U'' \sim \begin{cases} \lceil d/c \rceil + \operatorname{Bin}\Big(d - \lceil d/c \rceil,\, 1 - e^{-\lambda/c}\Big) & \text{with probability } 1-q, \\ 1 + \operatorname{Bin}\Big(d-1,\, 1 - e^{-\lambda/d}\Big) & \text{with probability } q, \end{cases} \tag{3.6}$$

where $q = \mathbf{P}\big[U_1' < \lceil d/c \rceil\big]$. Then $U'' \preceq U'$.

Proof. Writing $U' \mid E$ to mean U' conditioned on the event E, we claim that

$$U' \mid \{U_1' \geq \lceil d/c \rceil\} \succeq \lceil d/c \rceil + \mathrm{Bin}\Big(d - \lceil d/c \rceil,\, 1 - e^{-\lambda/c}\Big), \tag{3.7}$$

and

$$U' \mid \{U_1' < \lceil d/c \rceil\} \succeq 1 + \mathrm{Bin}\Big(d-1,\, 1 - e^{-\lambda/d}\Big). \tag{3.8}$$

The lemma then follows because conditional stochastic dominance implies stochastic dominance [SS07, Theorem 1.A.3, (d)].

Thus it just remains to confirm (3.7) and (3.8). Suppose $U_1' \geq \lceil d/c \rceil$. Then we are in Case 1, and $U' = \lceil d/c \rceil + U_2'$, where U_2' is the number of vertices in $\{u_1, \ldots, u_d\} \setminus \mathcal{V}$ visited by frogs returning from $\mathcal{V}$. Conditional on $\mathcal{V}$, the counts of frogs proceeding from $\mathcal{V}$ to each of $\{u_1, \ldots, u_d\} \setminus \mathcal{V}$ form a collection of independent $\mathrm{Poi}(\lambda \lceil d/c \rceil / d)$ random variables. Thus each vertex in $\{u_1, \ldots, u_d\} \setminus \mathcal{V}$ has an independent probability of $1 - e^{-\lambda \lceil d/c \rceil / d} \geq 1 - e^{-\lambda/c}$ of being visited by one of these frogs, showing that $U_2' \succeq \mathrm{Bin}(d - \lceil d/c \rceil, 1 - e^{-\lambda/c})$ and confirming (3.7).

Next, suppose that $U_1' < \lceil d/c \rceil$, and Case 2 is in effect. In this case, $U' = 1 + U_2'$, where U_2' is the number of vertices $u_2, \ldots, u_d$ visited by frogs returning from u_1. By the same reasoning as in the previous case, $U_2' \succeq \mathrm{Bin}(d-1, 1 - e^{-\lambda/d})$, confirming (3.8). □

Lemma 3.10. *Define*

$$h_{C,c} = h_{C,c}(\lambda, d) := \log\Big[\big(e^{-\frac{\lambda}{c} + \frac{\lambda}{d}} + 1 - e^{-\frac{\lambda}{c}}\big)^{d - \lceil d/c \rceil} + p\big(2 - e^{-\frac{\lambda}{d}}\big)^{d-1}\Big],$$

where p is the value defined in (3.5)*, which depends on C and c. We have*

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq \mathrm{Poi}\bigg(\lambda + \frac{\mu}{d+1} - h_{C,c}\bigg).$$

Proof. Combining (3.1) and $U'' \preceq \mathcal{U}\,\mathrm{Poi}(\lambda)$, it follows from [SS07, Theorem 1.A.3, (d)] that

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq \mathrm{Poi}\bigg(\frac{\mu}{d+1} + U''\frac{\lambda}{d}\bigg). \tag{3.9}$$

In light of Theorem 2, we would like to compute $-\log \mathbf{E} e^{-\frac{\lambda}{d}U''}$. Recalling the definition of U'' in (3.6), we use the fact that $\mathbf{E}x^{\mathrm{Bin}(n,p)} = (1 - p + px)^n$ to compute

$$\begin{aligned}\mathbf{E}e^{-\frac{\lambda}{d}U''} = (1-q)e^{-\frac{\lambda}{d}\lceil d/c \rceil}\Big(e^{-\frac{\lambda}{c}} + \big(1 - e^{-\frac{\lambda}{c}}\big)e^{-\frac{\lambda}{d}}\Big)^{d - \lceil d/c \rceil} \\ + qe^{-\frac{\lambda}{d}}\Big(e^{-\frac{\lambda}{d}} + \big(1 - e^{-\frac{\lambda}{d}}\big)e^{-\frac{\lambda}{d}}\Big)^{d-1}.\end{aligned}$$

Using the bound $q \leq p$ from Lemma 3.8 and the trivial bound $1 - q \leq 1$ in the first step, and factoring out $e^{-\lambda}$ in the second step,

$$\begin{aligned}\mathbf{E}e^{-\frac{\lambda}{d}U''} &\leq e^{-\frac{\lambda}{d}\lceil d/c \rceil}\Big(e^{-\frac{\lambda}{c}} + \big(1 - e^{-\frac{\lambda}{c}}\big)e^{-\frac{\lambda}{d}}\Big)^{d - \lceil d/c \rceil} \\ &\quad + pe^{-\frac{\lambda}{d}}\Big(e^{-\frac{\lambda}{d}} + \big(1 - e^{-\frac{\lambda}{d}}\big)e^{-\frac{\lambda}{d}}\Big)^{d-1} \\ &= e^{-\lambda}\Big[\big(e^{-\frac{\lambda}{c} + \frac{\lambda}{d}} + 1 - e^{-\frac{\lambda}{c}}\big)^{d - \lceil d/c \rceil} + p\big(2 - e^{-\frac{\lambda}{d}}\big)^{d-1}\Big].\end{aligned}$$

Thus,

$$-\log \mathbf{E}e^{-\frac{\lambda}{d}U''} = \lambda - h_{C,c}.$$

Using the above calculation and Theorem 2, we deduce that

$$\mathrm{Poi}\left(\frac{\mu}{d+1}+U''\frac{\lambda}{d}\right)\succeq \mathrm{Poi}\left(\lambda+\frac{\mu}{d+1}-h_{C,c}\right).$$

Together with (3.9), this completes the proof. □

Proof of Theorem 1. As we noted before, the lower bound is a consequence of [HJJ16a, Proposition 15], and we just need to establish the upper bound by showing that the frog model on $\mathbb{T}_d$ is almost surely recurrent with i.i.d.-$\mathrm{Poi}(2.28d)$ frogs per vertex for sufficiently large d. To apply our bootstrapping argument, we seek to show that for some $\delta>0$, it holds for all $\lambda\geq 0$ that $\mathcal{B}\,\mathrm{Poi}(\lambda)\succeq \mathrm{Poi}(\lambda+\delta)$. Considering the result of Lemma 3.10, we need to choose C and c such that $\mu/(d+1)-h_{C,c}(\lambda,d)>\delta$ for all $\lambda\geq 0$ and sufficiently large d. Recalling that $\mu=C(d+1)$, rearranging terms, and exponentiating both sides of the inequality, this is equivalent to showing that for some C, c, δ, and d_0 it holds that

$$\exp(h_{C,c}(\lambda,d))<e^{C-\delta}, \tag{3.10}$$

on the set $\{(\lambda,d)\colon \lambda\geq 0, d\geq d_0\}$.

Towards proving this, we start with the inequality

$$\begin{aligned}\exp(h_{C,c}(\lambda,d)) &= \left(e^{-\frac{\lambda}{c}+\frac{\lambda}{d}}+1-e^{-\frac{\lambda}{c}}\right)^{d-\lceil d/c\rceil}+p\left(2-e^{-\frac{\lambda}{d}}\right)^{d-1}\\ &\leq \left(1+e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}}-1\right)\right)^{d(1-\frac{1}{c})}+e^{-bd}2^{d-1}\end{aligned} \tag{3.11}$$

obtained by applying the bounds $2-e^{-\lambda/d}\leq 2$ and $d-\lceil d/c\rceil\leq d(1-1/c)$ and substituting the value of p from (3.5). Note that b depends on C and c. Now we bound each of the two terms on the right hand side of (3.11) for the right choice of C, c, and d_0.

Some calculus shows that for any d and c satisfying $d>c$, the first term is maximized in λ when $e^{\lambda/d}=d/(d-c)$. This demonstrates that if $d>c$, then

$$\begin{aligned}\left(1+e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}}-1\right)\right)^{d(1-\frac{1}{c})} &\leq \left(1+\left(\frac{d-c}{d}\right)^{d/c}\left(\frac{d}{d-c}-1\right)\right)^{d(1-\frac{1}{c})}\\ &= \left(1+\left(\frac{d-c}{d}\right)^{d/c}\frac{c}{d-c}\right)^{d(1-\frac{1}{c})}\\ &\leq \left(1+\frac{c}{d-c}\right)^{d(1-\frac{1}{c})}\\ &\leq \exp\left(\frac{d(c-1)}{d-c}\right).\end{aligned}$$

Thus, for any $\epsilon>0$, we can choose d_0 sufficiently large that for all $d\geq d_0$ and $\lambda>0$,

$$\left(1+e^{-\frac{\lambda}{c}}\left(e^{\frac{\lambda}{d}}-1\right)\right)^{d(1-\frac{1}{c})}\leq e^{c-1+\epsilon}. \tag{3.12}$$

The second term to be bounded, $e^{-bd}2^{d-1}$, vanishes as $d\to\infty$ when $b>\log 2$. Referring back to (3.5) and doing some algebra, we see that $b>\log 2$ when

$$C>-\log\left(1-\frac{1}{c}-\sqrt{\frac{\log 2}{2}}\right). \tag{3.13}$$

If this inequality holds, then for any $\epsilon>0$, we can choose d_0 large enough that for all $d\geq d_0$,

$$e^{-bd}2^{d-1}\leq\epsilon. \tag{3.14}$$

Finally, we set $c = 3.26$ and $C = 2.27$, which satisfies (3.13). Applying (3.12) and (3.14), for any $\epsilon > 0$, there exists d_0 such that for all $d \geq d_0$ and $\lambda \geq 0$,

$$\exp(h_{C,c}(\lambda, d)) \leq e^{2.26+\epsilon} + \epsilon.$$

Choosing $\epsilon, \delta > 0$ sufficiently small, we can bound this by $e^{2.27-\delta}$, confirming (3.10).

In all, we have shown that for $d \geq d_0$, if $\mu \geq 2.27(d+1)$ then for all $\lambda \geq 0$,

$$\mathcal{B}\,\mathrm{Poi}(\lambda) \succeq \mathrm{Poi}(\lambda + \delta). \tag{3.15}$$

Increasing d_0 as necessary, we can revise our assumption to $\mu \geq 2.28d$ for $d \geq d_0$. The rest of the proof is to use this to bootstrap our way to the conclusion that the number of visits to the root in the frog model is almost surely infinite given these assumptions, which proceeds identically as the last paragraph of Proposition 3. □

References

[AB05] David J. Aldous and Antar Bandyopadhyay, *A survey of max-type recursive distributional equations*, Ann. Appl. Probab. **15** (2005), no. 2, 1047–1110. MR-2134098

[AMP02] Oswaldo Alves, Fabio Machado, and Serguei Popov, *Phase transition for the frog model*, Electron. J. Probab. **7** (2002), no. 16, 1–21. MR-1943889

[DP14] Christian Döbler and Lorenz Pfeifroth, *Recurrence for the frog model with drift on* $\mathbb{Z}^d$, Electron. Commun. Probab. **19** (2014), no. 79, 13. MR-3283610

[GS09] Nina Gantert and Philipp Schmidt, *Recurrence for the frog model with drift on* $\mathbb{Z}$, Markov Process. Related Fields **15** (2009), no. 1, 51–58. MR-2509423

[HJJ16a] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *From transience to recurrence with Poisson tree frogs*, Ann. Appl. Probab. **26** (2016), no. 3, 1620–1635. MR-3513600

[HJJ16b] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *Recurrence and transience for the frog model on trees*, to appear in the *Annals of Probability*, available at arXiv:1404.6238, 2016.

[Hoe63] Wassily Hoeffding, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc. **58** (1963), 13–30. MR-0144363

[JJ16] Tobias Johnson and Matthew Junge, *Stochastic orders and the frog model*, available at arXiv:1602.04411, 2016.

[KZ16] Elena Kosygina and Martin P. W. Zerner, *A zero-one law for recurrence and transience of frog processes*, to appear in *Probability Theory and Related Fields*, available at arXiv:1508.01953, 2016.

[LP16] Russell Lyons and Yuval Peres, *Probability on trees and networks*, Cambridge University Press, 2016, Available at `http://pages.iu.edu/~rdlyons/`.

[MSH03] Neeraj Misra, Harshinder Singh, and E. James Harner, *Stochastic comparisons of Poisson and binomial random variables with their mixtures*, Statist. Probab. Lett. **65** (2003), no. 4, 279–290. MR-2039874

[Pop01] Serguei Yu. Popov, *Frogs in random environment*, J. Statist. Phys. **102** (2001), no. 1-2, 191–201. MR-1819703

[Ros16] Josh Rosenberg, *The frog model with drift on* $\mathbb{R}$, available at arXiv:1605.08414, 2016.

[SS07] Moshe Shaked and J. George Shanthikumar, *Stochastic orders*, Springer Series in Statistics, Springer, New York, 2007. MR-2265633

[TW99] András Telcs and Nicholas C. Wormald, *Branching and tree indexed random walks on fractals*, J. Appl. Probab. **36** (1999), no. 4, 999–1011. MR-1742145

[Yu09] Yaming Yu, *Stochastic ordering of exponential family distributions and their mixtures*, J. Appl. Probab. **46** (2009), no. 1, 244–254. MR-2508516